

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 429 201 A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 90311927.9

(51) Int. Cl.⁵: G06F 15/66

(22) Date of filing: 31.10.90

(30) Priority: 21.11.89 US 439859

(43) Date of publication of application:
29.05.91 Bulletin 91/22(34) Designated Contracting States:
DE FR NL(71) Applicant: **PICKER INTERNATIONAL, INC.**
595 Miner Road
Highland Heights Ohio 44143(US)(72) Inventor: **Vesel, John F.**
32216 Knobel Drive
Willowick Ohio 44095(US)
Inventor: **Mailey, Fred C.**

4121 Wilmington Road
S.Euclid, Ohio 44121(US)
Inventor: **Subramaniam, Pradipa**
443 Richmond Park West, Apt.505D
Richmond Heights, Ohio 44143(US)


Inventor: **Cecil, Robert A.**
33195 Wagon Wheel
Solon, Ohio 44139(US)

(74) Representative: **Pope, Michael Bertram**
Wingate
Central Patent Department Wembley Office
The General Electric Company, p.l.c. Hirst
Research Centre East Lane
Wembley Middlesex HA9 7PP(GB)

(54) **High efficiency image data transfer networks and methods.**

(57) A network (A) carries large image blocks among medical diagnostic equipment (20, 22, 24), archive computer (26), and data handling and display stations (28, 30). Four kilobyte packets of 4 megabyte image blocks are moved from transmit buffers (38) to queuing buffers (42). The order in which packets from the queuing buffers are transmitted on the network medium (10) is determined by a combination of an assigned priority, duration in the buffer, and a statistical availability of the addressed receiving node. A data link (52) at the receiving node includes an elasticity buffer (122) which stores a small plurality of bits, e.g. 5, to accommodate for variances in the clock speed of the transmitting and receiving nodes. A buffer table (60) monitors memory addresses at which preceding data packets corresponding to the same image are stored and provides address information to send each subsequently received packet. Before transmitting the image data packet, the transmitting node data link first sends a query to the receiving node to check availability to avoid the transmission of relatively large data packets to unavailable nodes. Each transmitting node normally passes a token after it has transmitted one packet of data. When the transmitting node holds the token for several sequential data packet

transmissions, it passes a restricted token which allows the other nodes to transmit high priority small messages that are much smaller than the image data packets.



EP 0 429 201 A2

HIGH EFFICIENCY IMAGE DATA TRANSFER NETWORKS AND METHODS

The invention relates to high efficiency image data transfer networks and methods. It finds particular application in conjunction with the transfer of image data among electronic medical diagnostic equipments such as computed tomography (CT) scanners, magnetic resonance imagers, digital x-ray equipments, and other equipments that generate, process or display digital image information. However, it is to be appreciated that the invention will also find application in other areas, particularly where large blocks of data are to be transferred.

Heretofore, medical diagnostic information has commonly been processed at each individual CT scanner, digital x-ray device, magnetic resonance imager or the like. Most often, the digital information is recorded on a removable disk or tape to facilitate carrying the image to other remote viewing locations. At the remote viewing locations, the radiologist could operate on the image data with various image enhancement algorithms and routines to enhance the diagnostic value for a particular application while the CT or other scanner was used to conduct further patient examinations. Rather than storing the image information on individual, portable disks, some hospitals have installed a central computer system with a large memory capacity for storing all digital images and other diagnostic data.

Difficulty has arisen in transferring the image data among the scanners, the central archiving computer, distributed viewing terminals, and other stations. An image is typically about 4 megabytes of data. With improvements in the scanners and in the processing speed of the generated data, images will be made up of progressively larger blocks of digital data. Blocks of data of this magnitude are most easily transmitted on dedicated, direct transmission or communication lines, that is, on communication lines which interconnect each viewing terminal, for example, with each potential source of images. However, such direct dedicated lines are expensive and switching among them is inefficient. Moreover, conflict management, such as the conflict between several viewing terminals all requesting images from the same archive computer or CT scanner at the same time, is difficult to resolve.

More commonly data is transferred among a plurality of stations or nodes on a common data bus. Various schemes have been developed for sharing the common bus, such as frequency or time division multiplexing. Most commonly, an asynchronous time division protocol is utilized.

To prevent packets of transmitted data from colliding or interacting on the data bus, various polling policies have been established. Commonly,

distributed polling policy is implemented, although central polling is also well known. Typical distributed or self polling schemes include token ring, logic ring, and the like.

On the whole, the hardware and protocols associated with central bus and ring architecture have high associated overhead. This overhead reduces the efficiency of the system and increases data transfer times. Moreover, these distributed systems and protocols were designed for applications with relatively small data packets by the standards of image data transfers. The prior art systems were also designed for flexibility to handle a wide range of applications. This flexibility again increases the overhead with no productive advantage in image data transfer applications.

Typically, at each data bus node, data is collected in the buffers of a file server from one or more sources. This data is recopied into a transport protocol hardware which adds its appropriate headers or other information. The data is copied from the buffers of the transport protocol to buffers of a queuing and multiplexing level, which again adds headers and other information to the data. The queuing and multiplexing level also oversees placing the data on the bus media in an efficient and fair way by the data link hardware. At the receive end, the data is again recopied from buffer to buffer as it is moved through the same layers, doubling the overhead.

It is an object of the present invention to provide new and improved data transfer networks and methods which overcome the above referenced problems and others.

According to a first aspect of the present invention there is provided a data transfer network comprising: a network media extending between each of a plurality of nodes; means for generating image representations connected with one of said nodes, each image representation including a relatively large block of image data; means for processing image data connected to another of said nodes; a first of said nodes including: buffer means for storing each of a plurality of image data packets, which taken together comprise an image data block, first data link means for transmitting each data packet on said network media, a second of the nodes including: second data link means for receiving the data packet from the network media; receiver memory means for storing at least one image data block, characterised in that: said first node includes a transmit data management means for controlling the transfer of data between said buffer means and said data link means and said second node includes receive data management

means for transferring image data packets from said second data link means to said receiver memory means, and, receiver memory address control means for monitoring addresses of said receiver memory means at which previously received image data packets are stored and providing receiver memory addresses for each received image data packet.

According to a second aspect of the present invention, there is provided a data transfer network comprising: a network media extending between each of a plurality of nodes; means for generating and processing large blocks of data connected with a plurality of said nodes; a first of said nodes including: data link means for receiving a plurality of data packets which taken together comprise one of the large blocks of data from the network media; and memory means for storing the large block of data; characterised in that said node further includes receive data management means for transferring each of the received data packets from said data link means to the memory means; and memory address control means for monitoring addresses of the memory means where previously received data packets are stored and for providing memory addresses for each subsequently received data packet.

According to a third aspect of the present invention, there is provided a data transfer network comprising: a network media extending between a plurality of nodes; means for processing large blocks of data connected with at least some of the nodes; a first of the nodes transmitting data to a second of the nodes, characterised in that said second node includes: an elasticity buffer into which data from the first node is clocked at a clocking rate of the first node and out of which the data is clocked at a clocking rate of the second node, said elasticity buffer having a sufficient storage capacity to accommodate variable numbers of bits of data attributable to differences in the clocking rates of the first and second nodes.

According to a fourth aspect of the present invention, there is provided a data transfer network comprising: a network media on which data is passed in n -bit format, where n is an integer, the network media extending between a plurality of nodes; means connected with said nodes for processing and handling blocks of m -bit format data, where m is an integer that is an even multiple of n ; at least one of said nodes including: data link means for receiving n -bit data from said network media, m/n serially connected n -bit buffers, clocking means for clocking n -bit data from said data link means serially through the n -bit buffers until said buffers are filled and, when said buffers are filled concurrently clocking said buffers to pass an m -bit format word to the associated one of the data

processing and handling means.

According to a fifth aspect of the present invention, there is provided a data transfer network comprising: a network media extending between each of a plurality of nodes; means for processing and handling data connected with the nodes; characterised in that a first of the nodes includes: means for storing a plurality of data packets, each packet being designated for transmission to one of the other receiving nodes and being assigned a relative transmission priority, each designated receiving node having a statistical probability of availability to receive a transmitted data packet; data link means operatively connected with said storing means and said network media for transmitting data packets serially on said network media; data packet order determining means for determining an order in which said data link means transmits the stored data packets on said network media in accordance with a weighted combination of the assigned priority of each packet and the statistical availability probability of the corresponding designated receiving node.

According to a sixth aspect of the present invention, there is provided a method of transferring data among a plurality of nodes interconnected by a network media, the method comprising: at a first of the nodes, storing at least one large block of image data, serially transmitting smaller data packets of the large data block on the network media to a second node and at a second node, serially receiving said data packets from the network media and storing each received packet in a memory; characterised in that said method includes the step of at said second node, keeping track of where packets corresponding to each data block being received packet by packet, are stored and directing each subsequently received packet to appropriate addresses in the memory.

According to a seventh aspect of the present invention, there is provided a method of transferring data among a plurality of nodes which are interconnected by a plurality of nodes, the nodes being clocked at clocking rates which are close to each other but not exactly the same, characterised in that said method comprises the steps of: transmitting data from a first of the nodes at a first clocking rate; at a second of the nodes, clocking the data into a buffer at the first clocking rate and clocking the data out of the buffer at a second clocking rate at which the second node is clocked.

According to an eighth aspect of the present invention, there is provided a method of transferring data between a plurality of devices that process data in m -bit format on a network media that transmits data in an n -bit format, characterised in that the method comprises the steps of: clocking n -bit data serially through a plurality of n -bit buffers:

when said n-bit buffers are filled, clocking said n-bit buffers concurrently to clock out an m-bit format word.

According to a ninth aspect of the present invention, there is provided a method of transferring packets of data among a plurality of nodes that are interconnected by a network media, the method comprising the steps of: at one of the nodes, storing a plurality of packets of data, each of which is waiting for transmission on the media, to a designated one of said plurality of nodes, each packet having a designated priority and each designated node having a statistical probability of availability to receive a packet; and characterised in that said method comprises the step of serially transmitting the awaiting packets on the network media in an order established by a combination of both the relative priority of each packet and the statistical availability of the corresponding designated node.

According to a tenth aspect of the present invention, there is provided a method of transferring large packets of data and small data transmissions among a plurality of nodes that are interconnected in a ring by a network media, characterised in that only a node which has a token is permitted to place a new data packet or a data transmission on the network media and characterised in that said method comprises the steps of: a first node, which has the token, intermittently transmitting new packets of data on the network media; between transmission of the data packets, said first node with the token passing a restricted token around the ring, said restricted token permitting a node possessing said restricted token to transmit smaller data transmissions over the network media but not the large data packets; as the restricted token is passed around the ring, a second node seizing the restricted token and transmitting a smaller data transmission to a third node.

An advantage of the present invention is that it more efficiently transfers image data or other large packets of data.

Another advantage of the present invention is that it is flexible and facilitates the adding of additional nodes to the network.

One image data transfer network and method in accordance with the present invention will now be described by way of example with reference to the accompanying drawings in which:-

Figure 1 is a system diagram of a data transfer network in accordance with the present invention;

Figure 2 is a diagrammatic illustration of an exemplary transmit and receive node of the network of Figure 1;

Figures 3A and 3B are a more detailed illustration of portions of the transmitting and receiving

nodes of Figure 2;

Figure 4 is a diagrammatic illustration of the data link assembly of Figures 2 and 3;

Figure 5 is a diagrammatic illustration of the unpacking means of Figure 4; and

Figure 6 is a schematic diagram of a packing scheme implemented by the packing means of Figure 4.

Referring to Figure 1, an image data transfer network A interconnects a plurality of stations or nodes B. In the preferred embodiment, the data transfer network is composed of a plurality of lengths of fiber optic cable 10 for transmitting data encoded light in a single direction from node to node. Of course, other transmission media may also be utilized as is known in the art. Each node retransmits the received information on the next length of fiber optic cable until the message has moved to each of the nodes and returned around the circle to the originating node.

Various nodes or stations may be provided. Typically, a plurality of scanners, such as a computed tomography (CT) scanner 20, a magnetic resonance imager 22, and a laser, digital x-ray, nuclear camera, or other digital imager 24 may be used. The imagers at these nodes non-invasively examine a patient and generate image data indicative thereof. The imagers may process the image data at each station and transmit completed images or may convey the raw image data along the network to another node which is interconnected with an appropriately programmed computer for processing the data. An archiving computer memory means 26 stores a large plurality of images. Because images typically contain about 4 megabytes of data, the archiving memory means may typically have a plurality of hard disks. Other nodes of the network are connected with operator control stations 28 from which the operator may control a scan, may view a selected image, operate on the selected image or images with various enhancement routines to improve their diagnostic value, and the like. The operator terminals may be interconnected with a hard copy apparatus for converting the electronically displayed diagnostic images into hard copy images on film, paper, or the like. Other nodes of the system may be connected with a more simple physician viewing station 30 at which a physician or radiologist may call up and view selected images for diagnostic purposes.

With reference to FIGURES 2 and 3A, each station node typically has both transmit and receive portions. When a request for an image or other data is made, a transmitting node portion associated with the request receiving station, e.g. one of the CT scanners 20, sends packets of a reconstructed diagnostic image representations from its image memory 32 to the requesting node. When

the reconstructed image representations are to be transferred by way of the network A, the image data is transferred through a data base means 34 to a file server 36. The file server typically receives and store an image in its transmit buffers 38. Although the images are referenced herein as being 4 megabytes, it is to be appreciated that images may be much larger or smaller.

Smaller packets of the image data are transferred through a transport protocol means 40 without recopying to queuing buffers 42 of a queuing and multiplexing means 44 under the control of a transmit data management means 46. The size of the packets is matched to the size of the queuing buffers. The queuing buffers are preferably as large as possible but are presently 4 kilobytes due to an unavailability of larger suitable buffer hardware.

The data in the queue buffers 42 may include high priority data in the priority buffer and lower priority data in a normal queue buffer. A transmit ordering means 48 controls the order in which the packets of data in queue buffers 42 are transferred.

In the preferred embodiment, the transmit order is based on three factors: the relative priorities of the data, the statistical probability that the designated receive node will be available to receive the data, and the duration that the data has been waiting in the queuing buffer. More specifically, each of the data packets is assigned a value V:

$$V = aP + bA + cT \quad (1),$$

where P is the relative priority of the data packets, A is the statistical availability of their designated receiving node, T is the duration or time that each of the data packets has been in the queuing memory, and a, b, and c are weighting factors. The weighting factors may be predetermined and assigned based on iterative adjustment, practical experience, or computer implemented system optimizing procedures. In the preferred embodiment, an adaptive control means 50 monitors the efficiency of the transmission of the data network and iteratively adjusts the weighting factors to improve and maintain the optimum efficiency with changing transmission patterns.

The data management means 46 adds appropriate header information to the data as it is moved from the file server buffers 38 to the queuing buffers 42. This header information includes a start delimiter, a destination node address, the source or transmitting node address, the 4 kilobyte data or packet field or information, an end delimiter, and a frame status.

When the transmitting node is authorized to transmit information, a data link means 52 receives the packet data from the designated queue buffer and the header information from the data management means 46 to formulate a message. The message is applied to the network media 10. Typically,

the data message moves to the next node in the unidirectional data transmission direction where it is received and retransmitted. This receive and retransmit procedure continues until the receiving node recognizes its address in the header.

with continuing reference to FIGURE 2 and further reference to FIGURE 3B, a data link means 52' at the receiving node retrieves the message and transmits it to a demultiplexing means 54. A receive data management means 56 receives the data packet and a decoding means 68 receives the multiplexing means applied headers. The header information is utilized to address a look-up table 60 which looks up the appropriate address for buffers 62 of a file server 64. The headers are used to determine to which of a potential plurality of images each data packet corresponds. For a given designated image, the data packets are received in order. Thus, each received 4 kilobyte field or packet of data is loaded into the receiving buffer memory means 62 in the locations immediately following the preceding data packet of the same image. The buffer table means 60 keeps track of the appropriate starting address for each received packet or the ending address for the precedingly received packet. The received data management means 56 transfers the data packet to the receive buffer memory means 62 at the addresses designated by the buffer table means 60. In this manner, the data is transferred directly to the receive buffer memories of the file server 64 without the inefficiency of recopying by the demultiplexing means 54 or a transport protocol means 66. A data base means 70 transfers the 4 megabyte images to a processor or other portions of the system 72 at the receiving node.

Various techniques may be utilized for adjusting the weighting factors of Equation 1. In a token ring protocol, the queries and the original messages return to the transmitting node with an indication whether they were accepted by the receiving node. A statistical table may be kept of the number of times the query or message comes back with an indication that each receiving node is busy. This information enables the transmit adaptive control means 50 to adjust the selection of packets from the buffers 42 to pick the one which has the best prospects for being received.

A multiplexing management system 74 receives information from a receive control buffer 76 concerning which image or other data is requested. The multiplex management means 74 may be connected with the data link means 52 for placing image data requests on the network A to be communicated to the station that has the requested information. These image data requests are received by the data link 52 of the transmitting buffer and communicated to the appropriate file server

and multiplexing means controls 78, 80. Alternatively, the requests for data can be transmitted by a separate data channel. The request is communicated to the file server transmit buffer control means 78 for controlling which image representation is to be moved in packets to the queue buffers 42 for transmission to the requesting station. Each information request includes priority information, which is conveyed to the multiplexor control 80 to provide the priority input regarding the order or transmission determination.

Various data management protocols may be utilized for managing the movement of data along the network A. In the preferred embodiment, a token ring type protocol is utilized in which messages are passed downstream from node to node around the ring. When a node detects its own address as the destination address, it copies the data packet into its buffers as described above. When the transmission returns to the transmitting node, the transmitting node strips the message from the ring and passes the token to the next node.

In the normal mode of operation, the token is passed after each transmission to assure equal access to the network by all nodes. The token is a control signal that is used to grant a node the privilege to transmit on the ring. A node that wishes to transmit captures the token and removes it from the ring. The transmitting node first sends a query to determine whether the destination node can accept the information packet. A positive acknowledgement is generally returned by the addressed receiving node in the form of trailing indicators that have been modified by the destination to designate its status. The transmitting node now formulates a message including headers as described above and a data packet - in the preferred embodiment, 4 kilobytes. The message is passed from node to node and copied by the destination node if it is available to receive. The destination node modifies the header information or trailing indicators to indicate whether or not it copied the information. When the message is returned to the transmitting node, it removes it from the network and notes whether or not the receiving node copied the information. If the information was copied, the transmitting node readies its next packet of information; if the information was not copied, the transmitting node prepares to retransmit the packet again. In either instance, the token is passed.

The protocol for controlling each node can be described as two separate state machines that communicate with each other - the receiver state and the transmitter state. Symbols within each packet arrive at the receiver asynchronously to the state of the transmitter. This requires that the receiver receive packets concurrently with other ac-

tions which may be being undertaken by the transmitter. The transmitter is informed of various receiver events, such as the reception of a claim packet. Notification occurs through signals that are passed from one state machine to the other. Global signals are used for events that cause responses from both machines. The receiver continuously monitors the signal stream that is being received and detects valid frames and tokens. packets whose destination address matches the nodes own address are copied if the buffer space is available and the appropriate process is notified. The transmitter is informed of tokens and frames that are received.

The receiver machine enters a listen state when it is initialized. This happens on power up or if a violation has occurred. Violations include a loss of synchronization with an upstream node. Receipt of a valid synchronization symbol indicates that synchronization has been reestablished. The receiver further scans the input symbol stream for commands that indicate the start of either the token or frame packets. The receiver also monitors for receipt of the token, the start of a frame, and the like. If data addressed to the receiver is copied into the buffers, an appropriate flag is set. This flag acknowledges the receipt to the transmitting node.

With reference to FIGURE 4, when the query comes back to the transmit node that the destination node is available to receive, it moves the data packet into a first in/first out (FIFO) transmit memory 100 of the data link 52. The data packets are set up for m-bit processors; whereas, the networks transfer n-bit data. Typically, $m = 32$ and $n = 8$. However, these values are based on hardware availability and are not constraints on the present invention. An n to m and m to n packing/unpacking means 102 unpacks the data from m bits to n bit format and packs the data from n to m bit format.

With reference to FIGURE 5, the unpacking means includes an m-bit buffer 104 and n m/n:1 multiplexors 106, e.g. eight 4:1 multiplexors. A clock means 108 clocks the m/n:1 multiplexors m/n times as fast the m-bit buffer 104. There are n one bit lines extending in parallel from the multiplexors to the transmit memory 100 to move n-bit packets of data into the transmit FIFO memory 100. After the multiplexors have been clocked m/n times, the m-bit buffer 104 is clocked to bring in the next m-bit word to be broken down into m/n n-bit words. The next m/n clock cycles strobe the multiplexing means 106 continuing the process.

With reference to FIGURE 6, when data is received, it is received in n-bit data segments and packed by the packing means into m-bit words. The packing means includes m/n n-bit buffers, in the illustrated embodiment four 8-bit buffers 110a, 110b, 110c, 110d. On the first clock pulse, the first

8-bits are loaded into the first 8-bit buffer 110a. On the second clock pulse, the first 8-bits are moved to the second buffer 110b and the next 8-bits are loaded into the first buffer 110a. Similarly, on the third clock pulse, the data is shifted one buffer such that three of them are filled. On the fourth clock pulse, the first 8-bits are shifted to the fourth buffer 110d, the next 8-bits into the third 8-bit buffer 110c, the second 8-bits into the second 8-bit buffer 110b and the last 8-bits into the first 8-bit buffer 110a. On the fifth clock pulse, all four 8-bit buffers are clocked simultaneously into a 32-bit buffer 112 forming the 32-bit long word. The number of buffers may be doubled to accommodate a 64-bit processor.

When the message is received at the downstream node, it is converted from an optical signal to an electrical transmission by a converter means 120 and loaded into an elasticity buffer means 122. The elasticity buffer is clocked by timing signals interspersed among the message to load information into the elasticity buffer. The elasticity buffer is also clocked by an internal clock 124 into a protocol defined communications layer 126. If the communications layer 126 recognizes its address in the message, it channels the information to a receive FIFO memory 128, through the packing means 102 and to the receive data management means 56 and the decoding means 58. Also, it clocks the message back to a transmitting means 130 which places the same message back on the network media 10.

It is to be appreciated that the clocks of the transmit and receive buffers transmit at the same frequency with an accuracy of four or five decimal places. However, clocks at the extremes of the tolerance range can build up a surplus or shortfall of several bytes in a 4 kilobyte transmission. Because the size of the packet is only limited by the size of the receive and transmit FIFO memories, larger packets may be transferred as larger FIFO memories become available. Accordingly, it is advantageous to size the elasticity buffer with sufficient capacity to assure that even with larger FIFO buffers and even when plural messages follow back to back or other extreme conditions occur, the elasticity buffer can accommodate a sufficient number of bytes to handle the full deviation in clock speed tolerances, for example, about 5 bytes. It is to be appreciated that the elasticity buffer has no address ports, data is just strobed in as it is received at the transmitting node clock rate and strobed out as it is ready for retransfer or copying at the receive node clock rate.

When the node seizes a token, it can generate packets on the ring. If the node needs to utilize the network for an extended duration, it issues a restricted token. Other nodes having high priority

small packets of data to transmit put out a priority flag. When a restricted token is passed around, the node with the priority flag can capture the restricted token and transmit a small packet of priority data. In this manner, associated equipment such as a laser imager which has a very small buffer, e.g. two lines, and must move its data around the network almost immediately to make room for the next lines, can keep control of the network while still allowing small data packets to be transferred in the time between lines.

Periodically, one of the nodes functions in a management capacity to verify the nodes in the network. It sends out a message with an address that is recognized by every node and instructions for each node to add its address to the end of the message. When the transmitting node again receives the message, it contains a list of each node on the network in order. The managing node then formulates an ordered node list which is transmitted to all nodes.

Each node continuously monitors the headers and trailing symbols of each message. The reception of frame indicators directs the nodes to monitor the fields which instruct each receiving node how to handle the message. The symbols can be indicative of the setting up of an ordered data node list, that information should be copied by all nodes, that the message may be stripped from the ring, whether the packet is valid, or the like.

A ring error recovery is accomplished with a claim process. When a node detects the absence of packet transmissions, it begins the claim process by generating a claim packet. Nodes receiving the claim packet with a source address less than their own pass on the received packets. A node that receives a claim packet with a source address greater than its own will remain in the claiming token state. A node that receives its own claim packet back will have won the arbitration and can generate a token. The failure of the node to receive a claim packet during the claim process indicates a serious ring fault. The node then enters a beacon process in which it generates a beacon packet that is receivable by all nodes. A node that receives its own beacon packet back having circumnavigated the ring, enters the claim process to begin initialization of the ring.

Claims

1. A data transfer network comprising: a network media (A, 10) extending between each of a plurality of nodes (B); means (20, 22, 24) for generating image representations connected with one of said nodes, each image representation including a relatively large block of image data; means for pro-

cessing image data connected to another of said nodes; a first of said nodes including: buffer means (38) for storing each of a plurality of image data packets, which taken together comprise an image data block, first data link means (52) for transmitting each data packet on said network media; a second of the nodes including: second data link means (52') for receiving the data packet from the network media (A,10); receiver memory means (62) for storing at least one image data block, characterised in that: said first node includes a transmit data management means (46) for controlling the transfer of data between said buffer means (38) and said data link means (52) and said second node includes receive data management means (56) for transferring image data packets from said second data link means to said receiver memory means (62), and, receiver memory address control means (60) for monitoring addresses of said receiver memory means (62) at which previously received image data packets are stored and providing receiver memory addresses for each received image data packet.

2. A network according to Claim 1 further including: receiver memory control means (74) which is responsive to operator input to request a selected image representation, the receiver memory control means (74) issuing an electronic request for the requested image representation; and the first node further including transmit control means (78) for controlling image representations to be transmitted, said transmit data management means (46) controlling the transfer of data packets of the selected image representation and said receiver control means (76) and said transmit control means (78) being operatively connected for the transfer of the electronic requests therebetween.

3. A network according to Claim 1 or 2 wherein the first node includes a plurality of queuing buffers (42), which hold image packets selected by different network nodes, a transfer of one of said data packets having a higher priority than the other, transmit throttling means (48) for controlling which of said data packets in said queuing buffers (42) is passed to said data link means (52) to be transmitted on said network media (A,10) first, said transmit throttling means (48) assessing the assigned priority of the data packets, the duration that each data packet has been in said queuing buffers (42), and a statistical probability that a designated receiving node is available to receive said data packet.

4. A network according to Claim 1, 2 or 3 wherein each data link includes an elasticity buffer (122) into which data is clocked at a clocking rate of data received from said network media (A,10) and out of which data is clocked at a clocking rate of the receiving data link (52'), the elasticity buffer (1-22) having sufficient storage to accommodate variable

numbers of bits of data attributable to differences in the clocking rates of said transmitted data and said receiver data link (52').

5. A network according to Claim 4 further including packing means (102) for packing data received from said elasticity buffer (122) in an n-bit format into an m-bit format, where m is greater than n, the packing means (102) including a plurality of n-bit buffers (110a, 110b, 110c, 110d), clocking means (124) for clocking incoming data serially through said n-bit buffers (110a, 110b, 110c, 110d) until the n-bit buffers are filled, said clocking means (124) concurrently clocking said n-bit buffers (110a, 110b, 110c, 110d) into an m-bit buffer (112) and clocking m-bit data from said m-bit buffer (112).

6. A network according to Claim 5 wherein said buffer means (122) stores data in m-bit format and further including an unpacking means (102) for converting the m-bit data to n-bit data, said unpacking means (102) including: an m-bit buffer (104) for receiving m-bit words; n multiplexor means (106) connected with said m-bit buffer (104); clocking means (108) for clocking said n multiplexor means (106) and said m-bit buffer (104), said clocking means (108) clocking the multiplexor (106) means about m/n times as fast as the m-bit buffer (104) such that each clocking of the n multiplexor (106) means generates an n-bit word for transmission on said network media (A,10).

7. A network according to any one of the preceding claims wherein each node includes management means for generating a receive query to a designated receiving node, when it is ready to transmit an image data packet, to determine whether said receive node is available to receive data packets and for transmitting the data packet with appropriate address headers when a reply to the query is received.

8. A network according to Claim 7 wherein a token permits the node possessing said token to transmit and the management means of the node with the token includes means for selectively generating a restricted token which is passed to other nodes to enable only small packets of data to be transmitted on the network media, the small packets of data being much smaller than the image data packets.

9. A network according to any one of the preceding claims further including packing means (102) for packing data received at the second data link means (52') in an n-bit format into an m-bit format, where a is an integer greater than n, said packing means (102) including m/n n-bit buffers (110a, 110b, 110c, 110d), clocking means for clocking incoming data serially through the n-bit buffers until the n-bit buffers are filled, said clocking means concurrently clocking the n-bit buffers into an m-bit buffer and clocking m-bit format data from the m-bit buffer.

10. A network according to any one of the preceding claims wherein the buffer means stores data in an m-bit format and data is transmitted on the network media (A, 10) in an n-bit format, where n is an integer and m is an integer that is an even multiple of n, and further including unpacking means (102) for converting the m-bit format to n-bit format data, the unpacking means including: an m-bit buffer (104) for receiving words of m-bit format; n multiplexor means (106) connected with the m-bit buffer; clocking means (108) for clocking said multiplexor means (106) and said m-bit buffer (104), said clocking means (108) clocking said multiplexor means (106) about m/n times as fast as said m-bit buffer (104) such that each clocking of said n multiplexor means (106) generates an n-bit format word for transmission on the network media (A, 10).

11. A data transfer network comprising: a network media (A,10) extending between each of a plurality of nodes (B); means (20, 22, 24, 26, 28) for generating and processing large blocks of data connected with a plurality of said nodes; a first of said nodes including: data link means (52') for receiving a plurality of data packets which taken together comprise one of the large blocks of data from the network media; and memory means (62) for storing the large block of data; characterised in that said node further includes receive data management means (56) for transferring each of the received data packets from said data link means (52') to the memory means (62); and, memory address control means (60) for monitoring addresses of the memory means (62) where previously received data packets are stored and for providing memory addresses for each subsequently received data packet.

12. A network according to Claim 11 wherein data link means (52') includes an elasticity buffer (122) into which said data packets are clocked in at a clocking rate of an upstream transmitting node and out of which data is clocked to said memory means (62) at a clocking rate of said first node.

13. A network according to Claim 11 or 12 wherein said memory means (62) stores data in an m-bit format and data is transmitted on the network media (A, 10) in an n-bit format and further including packing means (102) for packing n-bit format data received by said data link means (52') into m-bit format data for storage in said memory means (62).

14. A data transfer network comprising: a network media (A,10) extending between a plurality of nodes (B); means (32, 34, 36, 40, 44, 52, 52', 54, 66, 64, 70) for processing large blocks of data connected with at least some of the nodes; a first of the nodes transmitting data to a second of the nodes characterised in that said second node in-

cludes: an elasticity buffer (122) into which data from the first node is clocked at a clocking rate of the first node and out of which the data is clocked at a clocking rate of the second node, said elasticity buffer (122) having a sufficient storage capacity to accommodate variable numbers of bits of data attributable to differences in the clocking rates of the first and second nodes.

15. A network according to Claim 14 wherein data from the elasticity buffer (122) is in an n-bit format and which further includes packing means (102) for packing the n-bit format data into m-bit format data, where n is an integer and m is an integer which is an even multiple of n, said packing means (102) including: m/n serially connected n-bit buffers (110a, 110b, 110c, 110d) clocking means for clocking incoming data serially through said n-bit buffers (110a, 110b, 110c, 110d) until each is filled, the clocking means further concurrently clocking all m/n n-bit buffers (110a, 110b, 110c, 110d) to clock out an m-bit format word.

16. A data transfer network comprising: a network media (A,10) on which data is passed in n-bit format, where n is an integer, the network media extending between a plurality of nodes (B); means (32, 34, 36, 40, 44, 52, 52', 54, 64, 66, 70) connected with said nodes for processing and handling blocks of m-bit format data, where m is an integer that is an even multiple of n; at least one of said nodes including: data link means (52') for receiving n-bit data from said network media (A,10), m/n serially connected n-bit buffers (110a, 110b, 110c, 110d), until said buffers are filled and, when said buffers are filled, concurrently clocking said buffers (110a, 110b, 110c, 110d) to pass an m-bit format word to the associated one of the data processing and handling means.

17. A network according to Claim 16 wherein a second of the nodes includes: an m-bit buffer means (104) for storing a plurality of packets of m-bit data to be transferred on the network media, each stored packet awaiting transmission to a designated other receiving node, each of the packets having a relative priority and each of the designated receiving nodes having a statistical probability of availability to receive packets; unpacking means (102) for converting data packets stored in m-bit buffer means (104) to n-bit format data for transmission on the network media (A,10); a second data link (52) for transmitting the n-bit data on the network media (A,10); and, means for selecting an order in which the data packets are passed from the m-bit buffer means (104) to said unpacking means (102) and said data link means (52) in accordance with a combination of the relative priority of each packet and the statistical availability of the designated receiving node for each packet.

18. A data transfer network comprising: a network

(A,10) media extending between each of a plurality of nodes (B); means for processing and handling data connected with the nodes (B); means for processing and handling data connected with the nodes (B); characterised in that a first of the nodes (B) includes: means (38) for storing a plurality of data packets, each packet being designated for transmission to one of the other receiving nodes and being assigned a relative transmission priority, each designated receiving node having a statistical probability of availability to receive a transmitted data packet; data link means (52) operatively connected with said storing means (38) and said network media (A,10) for transmitting data packets serially on said network media (A,10); data packet order determining means for determining an order in which said data link means (52) transmits the stored data packets on said network media (A,10) in accordance with a weighted combination of the assigned priority of each packet and the statistical availability probability of the corresponding designated receiving node.

19. A method of transferring data among a plurality of nodes (B) interconnected by a network media (A,10), the method comprising the steps of: at a first of the nodes, storing at least one large block of image data, serially transmitting smaller data packets of the large data block on the network media (A,10) to a second node and at a second node, serially receiving said data packets from the network media and storing each received packet in a memory (62); characterised in that said method includes the step of at said second node, keeping track of where packets, corresponding to each data block being received packet by packet, are stored and directing each subsequently received packet to appropriate addresses in the memory (62).

20. A method of transferring data among a plurality of nodes which are interconnected by a plurality of nodes, the nodes being clocked at clocking rates which are close to each other but not exactly the same, characterised in that said method comprises the steps of transmitting data from a first of the nodes at a first clocking rate; at a second of the nodes, clocking the data into a buffer at the first clocking rate and clocking the data out of the buffer at a second clocking rate at which the second node is clocked.

21. A method of transferring data between a plurality of devices that process data in m-bit format on a network media that transmits data in an n-bit format, characterised in that the method comprises the steps of: clocking n-bit data serially through a plurality of n-bit buffers (110a, 110b, 110c, 110d); when said n-bit buffers (110a, 110b, 110c, 110d) are filled, clocking said n-bit buffers (110a, 110b, 110c, 110d) concurrently to clock out an m-bit format word.

22. A method of transferring packets of data among a plurality of nodes (B) that are interconnected by a network media, the method comprising the steps of: at one of the nodes, storing a plurality of packets of data, each of which is waiting for transmission on the media (A,10), to a designated one of said plurality of nodes, each packet having a designated priority and each designated node having a statistical probability of availability to receive a packet; and characterised in that said method comprises the step of serially transmitting the awaiting packets on the network media (A,10) in an order established by a combination of both the relative priority of each packet and the statistical availability of the corresponding designated node.

23. A method according to Claim 22 further including before each packet is transmitted the step of transmitting a query to the designated receiving node to inquire whether it is available to receive a data packet.

24. A method of transferring large packets of data and small data transmissions among a plurality of nodes (B) that are interconnected in a ring by a network media (A,10) characterised in that only a node which has a token is permitted on the network media and characterised in that said method comprises the steps of: a first node, which has the token intermittently transmitting new packets of data on the network media (A,10); between transmission of the data packets, said first node with the token passing a restricted token around the ring, said restricted token permitting a node possessing said restricted token to transmit smaller data transmissions over the network media (A,10) but not the large data packets; as the restricted token is passed around the ring, a second node seizing the restricted token and transmitting a smaller data transmission to a third node.

25. A method according to Claim 24 wherein each data packet transmitted on the network media (A,10) between an originating node and a destination node is received and retransmitted by each intervening node, the packet being clocked into each intervening node at a clock rate of an immediately upstream node and being clocked out of the intervening node at a clock rate of the intervening node.




FIG. 1




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4

